Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

An Adaptive Framework for Robust Energy Forecasting under Concept Drift and Feature Uncertainty

Francesco Marcato[a], Alessio Santecchia[a], Manuel Ruivo de Oliveira[a], Francesco Silvestri[b], Rafael Castro-Amoedo[c]*

[a] Emissium Labs Unipessoal LDA, Alcácer do Sal, Portugal
[b] AIDA Lab (Algorithms for Intelligent Data Analytics), Department of Information Engineering, University of Padova
[c] Centro de Recursos Naturais e Ambiente (CERENA), Instituto Superior Técnico, Lisbon, Portugal;
* Corresponding Author: rafael.amoedo@tecnico.ulisboa.pt

ABSTRACT

The rapid integration of renewable energy sources is increasing the volatility and non-stationarity of modern power systems, posing significant challenges for data-driven forecasting models. In particular, concept drift and uncertainty in exogenous inputs such as weather forecasts can severely degrade predictive performance over time. This work proposes a lightweight two-layer forecasting framework that decouples prediction from adaptation. A traditional offline regression model is augmented by an online meta-learner that continuously generates adaptive meta-features, enabling the system to respond to structural changes and noisy inputs without repeated retraining. The framework is evaluated on two real-world case studies. First, concept drift is addressed in nuclear power production forecasting, where abrupt and gradual capacity changes are inferred through an online meta-learner. Second, feature uncertainty is mitigated in day-ahead solar production forecasting by correcting noisy weather forecast inputs. Across both scenarios, the proposed approach consistently outperforms single-layer baselines, reducing root mean squared error by up to 10% and maintaining robust performance over multi-year horizons without retraining. The results demonstrate that meta-learning provides a practical and computationally efficient mechanism for improving forecast robustness in non-stationary energy systems, with applicability to a wide range of power-system forecasting problems.

Keywords: Algorithms, Big Data, Modelling, Artificial Intelligence, Energy

INTRODUCTION

The increasing penetration of renewable energy sources is fundamentally changing the operating conditions of modern power systems. Weather-dependent generation, evolving market mechanisms, and policy-driven structural changes are introducing significant non-stationarity into electricity production and demand time series. Accurate forecasting remains essential for grid operation, market participation, and emerging carbon-aware applications [1, 2], yet conventional data-driven models often struggle when historical patterns no longer hold.

Supervised learning models trained offline remain widely used in energy forecasting applications [3]. While advanced deep learning architectures and sliding-window gradient boosting ensembles represent the current state-of-the-art for raw predictive accuracy, their high computational overhead and requirement for continuous retraining limit their scalability in real-time operational environments. However, these models implicitly assume stationarity in the underlying data-generating process. Violations of this assumption lead to systematic degradation of predictive performance over time, motivating the development of adaptive forecasting strategies.

Concept Drift and Adaptive Forecasting

Concept drift refers to changes in the statistical relationship between input variables and target outputs over time as described by Tsymbal [4], who formalised the concept and identified several drift patterns, including abrupt, gradual, and incremental drift. In energy systems, concept drift arises from structural events such as power plant commissioning or decommissioning, fuel

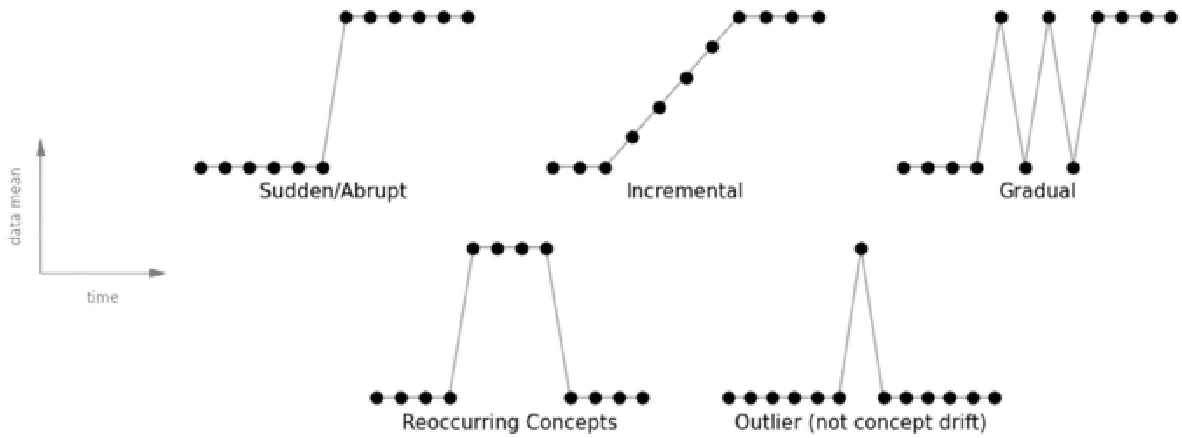

Figure 1. – Canonical drift patterns

switching, and regulatory interventions.

These drift types are illustrated schematically in **Figure 1**, which provides a conceptual reference for the experimental scenarios studied later in this work, and illustrates typical concept drift patterns, motivating the need for adaptive forecasting strategies capable of handling both abrupt and gradual changes.

Several approaches have been proposed to address concept drift in regression problems, including sliding-window retraining [5], weighted learning [6], and fully online adaptation [7]. While effective, retraining-based methods are computationally expensive, fully online learners may suffer from instability. On the other hand, Ensemble-based adaptive methods offer a compromise by combining robustness and adaptability.

Adaptive ensemble methods have gained significant attention for non-stationary data streams. Gomes et al. [8] introduced the Adaptive Random Forest (ARF), which integrates online bagging, per-tree drift detection, and background learners to enable localised adaptation. ARF has been successfully applied to a variety of streaming regression problems [9].

Efficient open-source implementations of these methods are provided by modern streaming machine-learning libraries, such as the River library in Python [10], facilitating their deployment in operational forecasting systems.

Uncertainty in Exogenous Inputs

In addition to structural drift, which refers to physical changes in the system (like equipment aging) rather than the statistical shifts of concept drift, forecasting accuracy is strongly affected by uncertainty in exogenous inputs. Renewable generation forecasting, particularly for solar and wind power, relies heavily on meteorological predictions that are inherently uncertain. Errors and biases in weather forecasts propagate directly into production forecasts, often resulting in large prediction errors [11].

Various strategies have been proposed to mitigate this effect, including probabilistic forecasting [12], noise injection during training [13], and ensemble-based uncertainty modelling [14]. More recently, meta-learning approaches have been explored to correct biased, noisy, or missing inputs online without retraining the main forecasting model [15].

Contribution

This paper proposes a two-layer meta-learning framework designed to address both concept drift and feature uncertainty without repeated retraining of the main forecasting model. The key contributions are:

1. A unified two-layer architecture that separates prediction and adaptation through meta-feature generation.
2. An online meta-learning strategy that adapts to structural changes and noisy inputs in streaming data.
3. Validation on real-world energy datasets covering both capacity-driven concept drift and weather-driven feature uncertainty.

METHODOLOGY

The proposed framework combines an online meta-learner with a base forecaster that undergoes offline training. In the present implementation, the base forecaster is a linear regression model trained using ordinary least squares (implemented using river [0.23.0], scikit-learn [1.7.2]). All experiments were conducted in Python [3.12.10]. The model is trained once on historical data, and its coefficients remain fixed during the evaluation

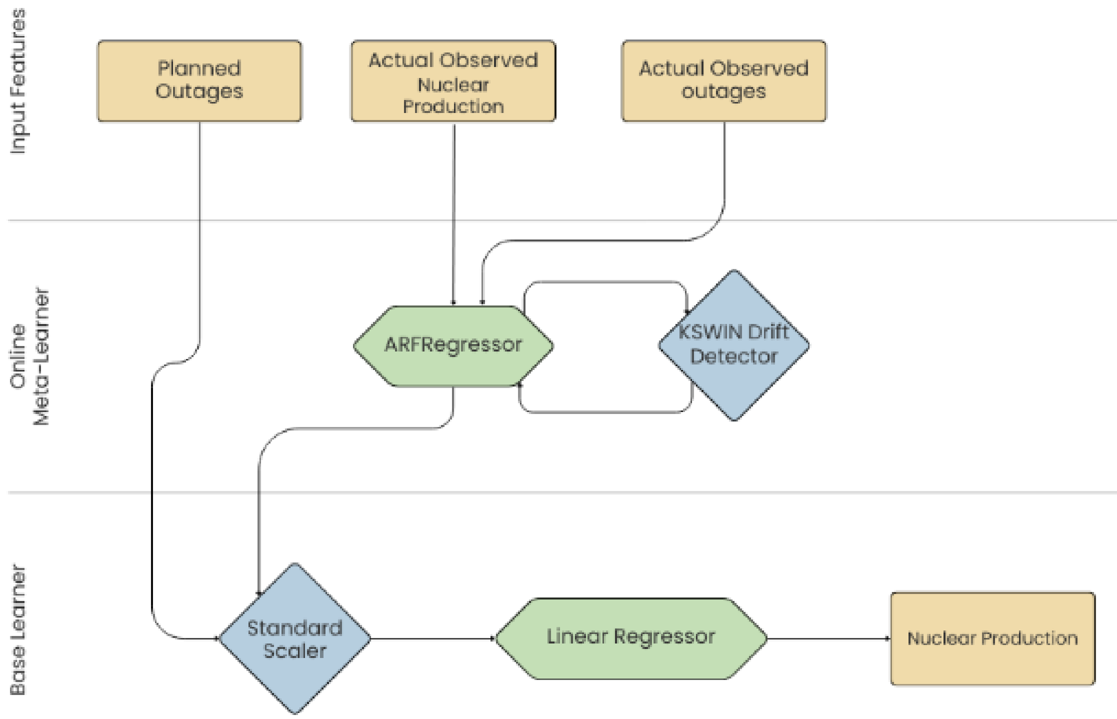

Figure 2. – Graphical representation of the two-layer framework, with an online meta-learner estimating operational capacity, which is appended to the offline forecaster as a meta-feature

period. The choice of a simple regression model for the base forecaster was deliberate. While modern forecasting often employs complex ensembles, utilizing a static offline baseline allows the isolated impact of the adaptive meta-feature to be measured cleanly, proving that our lightweight layer can rescue legacy models without the need for computationally expensive sliding-window retraining. This designation signifies that the model is pretrained on historical data and its internal weights remain fixed throughout the deployment phase. To enable adaptation without retraining the forecaster, the system augments the input with meta-features generated online.

Predictions $\hat{y}_t$ at time t are computed as:

$$\hat{y}_t = f(x_t, m_t),$$

where $x_t$denotes the original feature vector and m_t is a meta-feature generated by:

$$m_t = g(z_t),$$

with $g(\cdot)$ representing the meta-learner and z_t auxiliary information such as recent observations, forecast errors, or uncertain inputs. During deployment, the framework operates sequentially on the data stream. At each time step, the meta-learner receives the most recent observations and updates its internal state. The resulting meta-feature is appended to the original feature vector and passed to the fixed offline forecaster, which produces the final prediction. This streaming workflow enables continuous adaptation without modifying the parameters of the base forecasting model.

The overall architecture of the proposed two-layer framework is shown in **Figure 2**, which illustrates the interaction between the fixed offline forecaster and the online meta-learner. The separation of prediction and adaptation is central to the proposed approach and enables efficient handling of non-stationarity. In the nuclear forecasting case study (see below), the meta-learner estimates a latent variable representing the effective operational capacity of the nuclear fleet. This estimate is derived incrementally from the production time series and the outage signal available in the data stream. At each time step, the Adaptive Random Forest updates its estimate using the newly observed production value, producing a capacity-related feature that reflects the current operational state of the system. This estimate is then appended to the original feature vector as a single additional input to the offline forecaster.

The meta-learner is implemented using an Adaptive Random Forest operating in streaming mode as described by Gomes et al. [7]. The Adaptive Random Forest implementation used in this work is based on the River library [10]. Across both case studies the ARFRegressor always uses the default (n_models=10) unless otherwise stated. In the first case study, we override the default detector with KSWIN, whereas in the second case study, we use the default ARF configuration, which includes ADWIN detectors. The detectors act to identify distributional changes and trigger local model updates. Unless otherwise specified, the remaining hyperparameters follow the default configuration provided in the River implementation. Internal drift detection mechanisms allow individual trees within the ensemble to adapt locally, while the forecaster remains fixed, reducing computational overhead compared to retraining-based approaches. Models are compared using a set of metrics, including the coefficient of determination (R^2), the Root Mean Squared Error (RMSE), the Mean Absolute Error (MAE), the Mean Squared Error (MSE) and the Symmetric Mean Absolute Percentage Error (sMAPE).

All experiments were implemented in Python using the River streaming machine learning library (version 0.23.0). To ensure exact reproducibility of the Adaptive Random Forest's stochastic processes, the global random seed was fixed at 42 across all experiments. The full experimental workflow, including data preprocessing and model configuration, follows the methodology described above and can be reproduced using the same datasets and parameter settings.

Case Study I: Concept Drift in Nuclear Power Forecasting

The first case study evaluates the framework under structural concept drift in nuclear power production. Nuclear generation is generally stable but subject to abrupt changes due to plant shutdowns, making it a suitable testbed for drift handling.

The meta-learner is tasked with estimating the operational capacity of the nuclear fleet, a latent variable not directly observable from production data alone. This estimated capacity is provided as a meta-feature to an offline regression model trained prior to major structural changes. Experiments were conducted on national-level nuclear production data for Germany, Switzerland, and the Netherlands, sourced from the ENTSO-E transparency platform [17]. The dataset consists of hourly aggregated production measurements spanning the period 2018–2024. Production values were aligned with outage information obtained from the European Energy Exchange (EEX). The resulting time series was cleaned for missing values and resampled to hourly resolution to ensure temporal consistency across all variables.

Case Study II: Feature Uncertainty in Solar Power Forecasting

The second case study addresses feature uncertainty in day-ahead solar power forecasting. Solar generation forecasts rely heavily on weather predictions, which are inherently uncertain and may exhibit systematic bias.

In this scenario, the meta-learner operates as a feature uncertainty corrector. It learns the relationship between forecasted and observed solar radiation in an online fashion and outputs corrected radiation estimates as meta-features. The forecaster uses these corrected inputs to predict solar production.

The framework was evaluated on aggregated Swiss solar production data using three configurations: a baseline model trained on forecasted weather (using data from 2023-01-01 to 2024-01-01), the proposed two-layer model with corrected features, and an upper-bound model trained on observed weather over the same training interval. Solar radiation forecasts were obtained from the Visual Crossing Weather API and correspond to day-ahead meteorological predictions. Observed solar radiation values were aligned with the forecasts at hourly resolution. The forecasting model uses the radiation forecast together with the meta-corrected radiation estimate produced by the meta-learner.

RESULTS AND DISCUSSION

Structural concept drift: Nuclear Power Production

Nuclear power production constitutes a particularly challenging forecasting problem under concept drift. Unlike renewable generation, which is primarily driven by stochastic exogenous variables, nuclear production is dominated by structural constraints, namely installed capacity and plant availability. When these constraints change due to policy decisions or outages, the mapping between historical inputs and future outputs is fundamentally altered.

This effect is particularly evident in Germany, where the progressive nuclear phase-out resulted in multiple abrupt capacity reductions over the study period. The timing and magnitude of these events are shown in **Figure 3**, which highlights the strongly non-stationary nature of the production time series. Switzerland and the Netherlands exhibit milder forms of drift, characterised by gradual capacity changes and shorter outage periods.

These three countries, therefore, provide a spectrum of drift regimes, allowing the robustness of the proposed framework to be assessed under increasing levels of structural disruption.

We compared our proposed Two-Layer Framework (TLF) against a Single-Layer (SL) baseline defined as an offline linear regression model using only the outage signal as input. In the TLF, an online Adaptive Random Forest meta-learner estimates a latent capacity-related feature from the data stream (hourly resolution), and a second, offline linear regression main model predicts production using both outage and the meta-learner's output. Models were trained on 2018-01-01 to 2020-01-01 and evaluated on 2020-01-01 to 2024-01-01. Forecasts were generated sequentially in a walk-forward evaluation setting. At each time step in the test period, the meta-

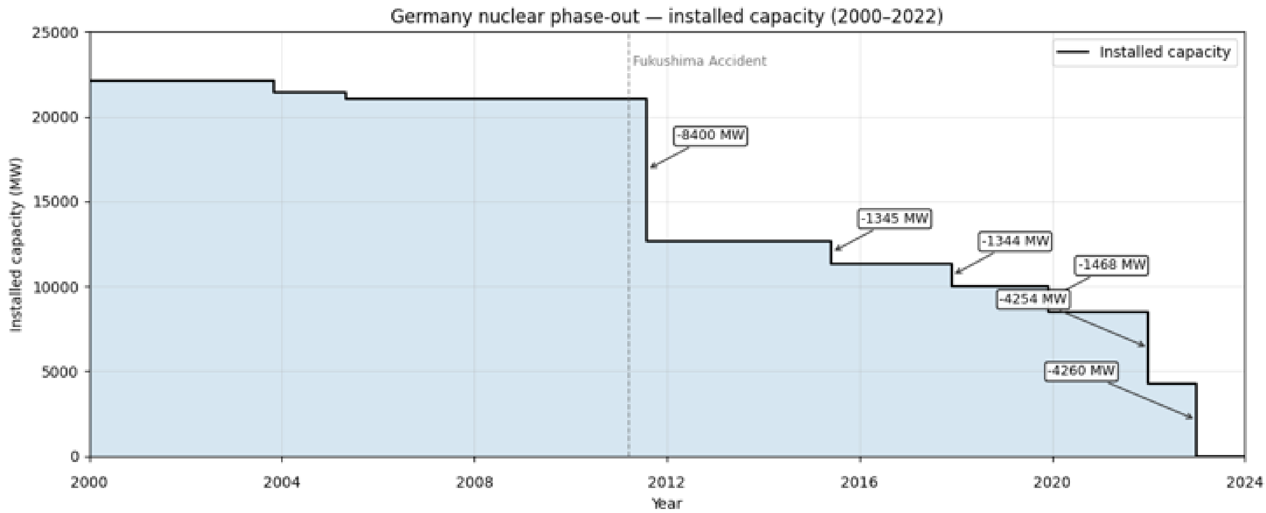

Figure 3. – Timeline for Germany's nuclear power plants phase-out culminating in April 2023. Source : Mez et al. [16].

learner was updated using the newly observed production value before the next prediction was produced.

Model performance was evaluated using several complementary error metrics, including RMSE, MAE, sMAPE, and the coefficient of determination (R^2). These metrics capture both absolute prediction accuracy and relative forecasting performance across different production scales. A quantitative comparison of forecasting performance is reported in **Table 1**, which aggregates RMSE, MAE, sMAPE and R^2 metrics for Germany, Switzerland, and the Netherlands.

Table 1 – Performance comparison of a single-layer (SL) traditional model and our two-layer framework with meta-learner (TLF) for Germany (DE), Switzerland (CH) and the Netherlands (NL).

Country	Model	R^2 [-]	RMSE [MW]	MAE [MW]	sMAPE [%]
DE	**SL**	-1.5547	4525	3714	71.8
	TLF	0.9824	375	293	40.0
CH	**SL**	0.6966	369	348	14.4
	TLF	0.9732	110	34	1.8
NL	**SL**	0.9644	26	10	12.3
	TLF	0.9698	24	6	9.5

Table 1 shows that the two-layer framework consistently outperforms the single-layer baseline in the presence of concept drift. In the Netherlands, where structural changes are limited, both models achieve similar performance, indicating that the meta-learning layer does not degrade accuracy under quasi-stationary conditions. This is an important result, as it demonstrates that the proposed approach introduces no penalty when adaptation is not strictly required.

In Switzerland, where gradual capacity changes and recurring outages occur, the two-layer framework yields a clear improvement over the baseline. RMSE reductions on the order of several percentage points are observed, alongside a noticeable increase in R^2. This indicates that the meta-learner is able to track slow structural changes effectively, providing the forecaster with a more accurate representation of the system's operational state.

Germany represents the most challenging case. Here, the baseline model rapidly loses predictive power following major shutdown events, as it continues to rely on outdated capacity assumptions. As shown in **Table 1**, the baseline model for Germany yields a severely negative R^2 (-1.5547). It is important to note that the coefficient of determination (R^2) measures the model goodness-of-fit and is not strictly equivalent to the squared Pearson correlation coefficient (r^2). Consequently, R^2 is not bounded at zero and can drop below -1 when a model's predictions perform substantially worse than simply predicting the historical mean of the observed data. In this case, because the static baseline model continuously relies on outdated capacity assumptions after abrupt phase-out events, its predictions exhibit extreme systematic bias, mathematically resulting in this highly negative value. The two-layer framework significantly mitigates this degradation, achieving substantially lower error metrics over extended periods. While performance still deteriorates during extreme shutdown phases, the meta-learning approach preserves a meaningful level of predictive accuracy where the baseline effectively fails.

While aggregate metrics provide a useful summary,

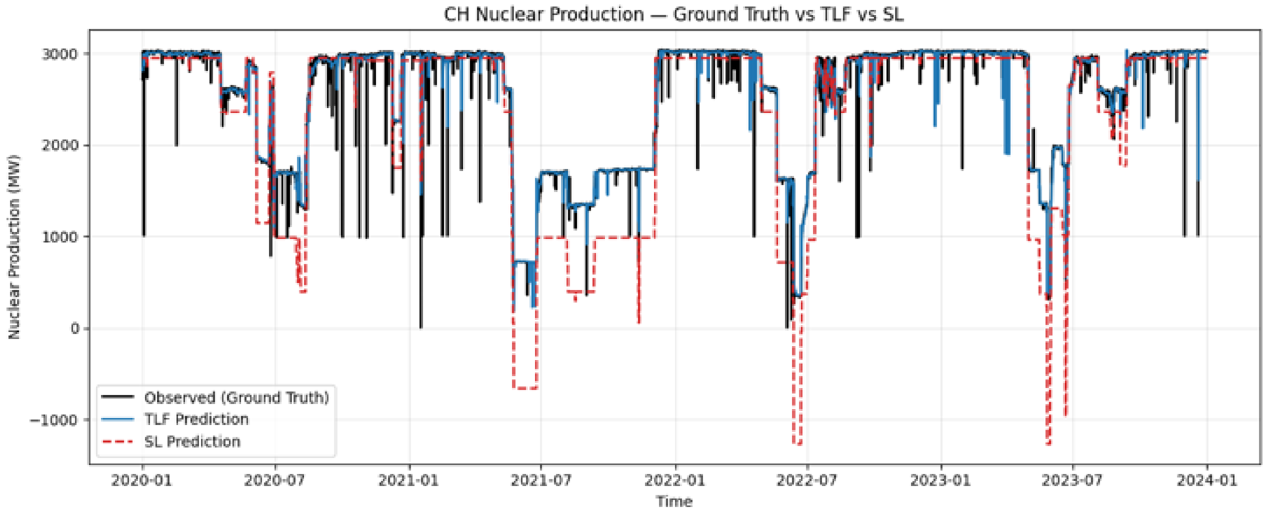

Figure 4. – Predictions of the Swiss nuclear production using the two-layer framework. The model tracks structural change better than the single-layer baseline.

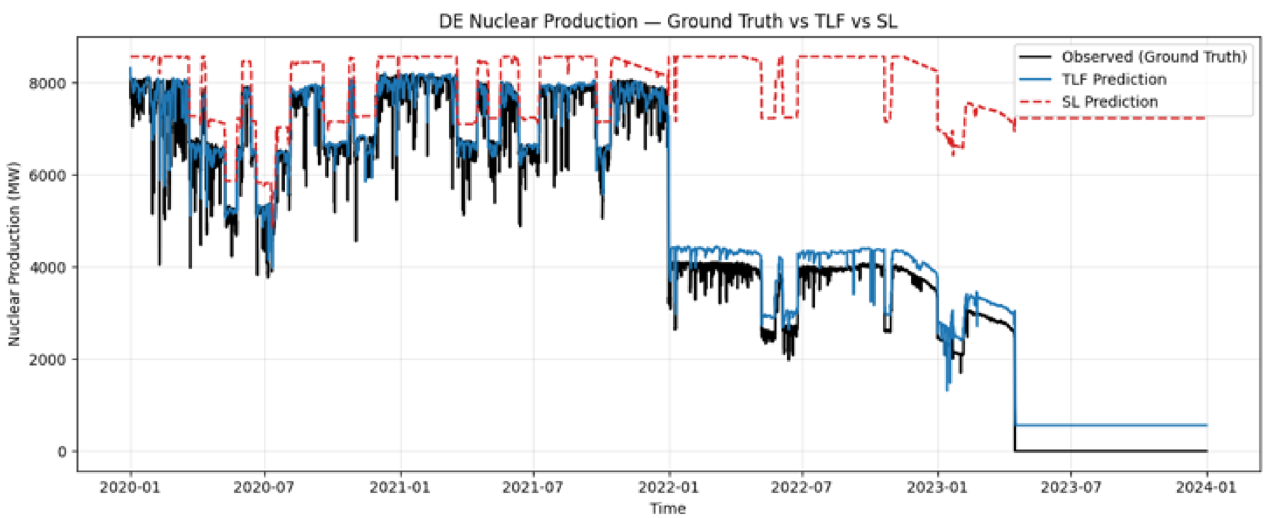

Figure 5. – Predictions of the German nuclear production using the two-layer framework. Despite improvements over the baseline, large abrupt capacity shutdowns remain challenging.

they obscure the temporal dynamics of adaptation. These dynamics are illustrated in **Figures 4 and 5**, which show predicted and observed nuclear production for Switzerland and Germany, respectively.

Figure 4 demonstrates how the two-layer framework maintains close alignment with observed production over time, even as capacity evolves. In contrast, the baseline model exhibits systematic bias during extended outage periods, reflecting its inability to adapt to changing structural constraints.

Figure 5 provides a more nuanced picture. Following abrupt shutdown events, both models experience an immediate increase in error. However, the two-layer framework adapts more rapidly, reducing bias and variance over subsequent time steps. This behaviour reflects the ability of the Adaptive Random Forest meta-learner to detect distributional change and update its internal representation accordingly.

Importantly, the results also reveal the **limits of adaptation**. When nuclear capacity approaches zero, the latent-variable assumption underpinning the meta-learner becomes invalid. In such cases, the model cannot fully recover, highlighting a fundamental limitation of any approach that relies on historical production patterns. This observation is critical and underscores the importance of explicitly acknowledging failure modes in adaptive forecasting systems.

From an engineering perspective, these results demonstrate that the proposed framework offers a practical alternative to frequent retraining. Structural capacity

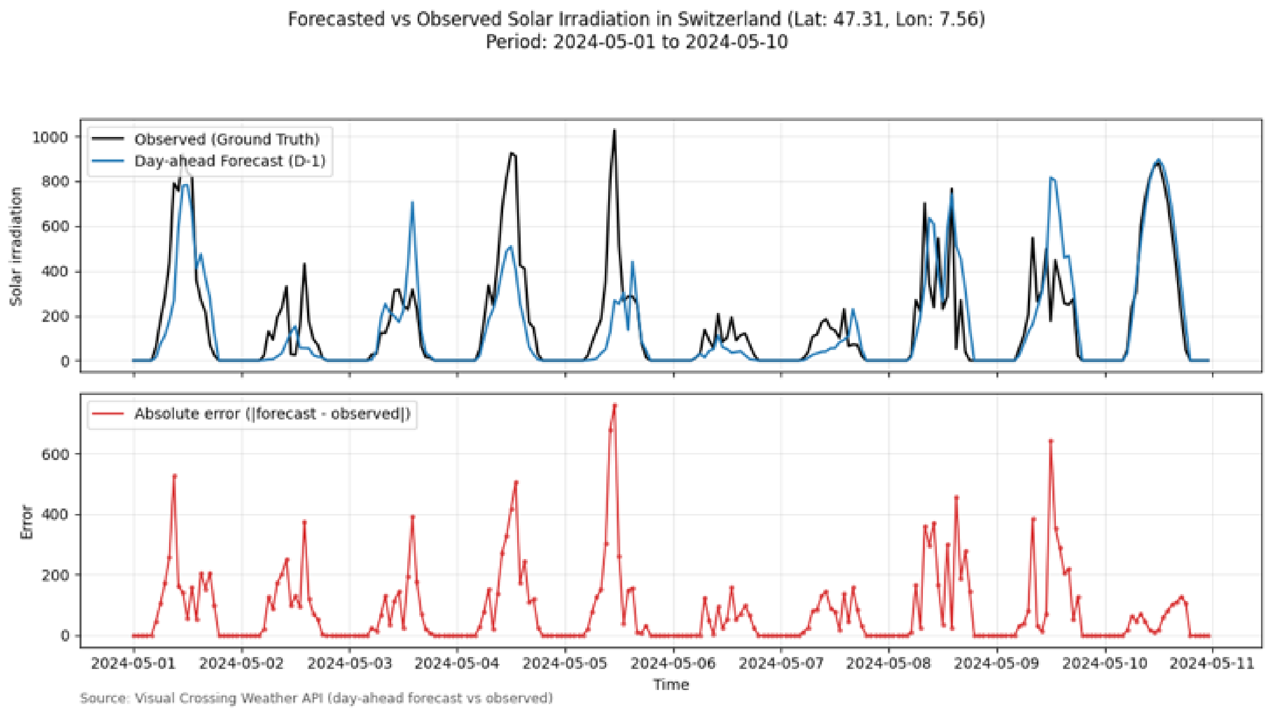

Figure 6. – Forecasted vs. observed solar irradiation in Switzerland (Lat: 47.31, Lon: 7.56) over the period May 1 to May 10, 2024. Data source: Visual Crossing Weather API. The top panel shows day-ahead forecasts (blue) deviating from observations (black), while the bottom panel tracks the absolute error, demonstrating recurrent inaccuracies.

changes are handled internally by the meta-learner, reducing the need for manual intervention or model redevelopment. This is particularly relevant for operational forecasting systems managing multiple assets or regions simultaneously.

Moreover, the results suggest that the two-layer approach is especially well suited to **moderate-to-strong drift regimes**, where structural changes occur but do not completely invalidate the historical relationship between inputs and outputs.

Feature uncertainty: Solar Power Production

In contrast to nuclear power, solar power production is dominated by uncertainty in exogenous inputs, particularly weather forecasts. Even when the underlying physical relationship between irradiance and production remains stable, errors in meteorological predictions introduce noise and bias that directly affect forecast accuracy.

This effect is illustrated in **Figure 6**, which compares forecasted and observed solar radiation. The figure reveals both random dispersion and systematic deviations, indicating that forecast errors are neither negligible nor purely stochastic.

Forecasting performance is summarized in **Table 2**. We compare the Baseline model against our Two-Layer Framework, as well as an "Upper Bound" (Oracle) model. The Upper Bound model was trained and tested using observed (perfect) weather data rather than forecasts; it represents the theoretical limit of performance achievable if meteorological predictions were error-free.

Table 2 – Performance comparison of the Baseline, Two-Layer, and Upper Bound models for Swiss solar production forecasting (Test set: 01.01.24 to 01.01.25). Metrics are computed on hourly aggregated data.

Model	**R^2 [-]**	**RMSE [MW]**	**MAE [MW]**
Baseline	0.7183	436.52	215.84
Two-Layer	0.7711	393.50	198.43
Upperbound	0.7740	391.34	185.34

Table 2 shows that the baseline model suffers from elevated RMSE and MAE due to error propagation from uncertain inputs. Introducing the meta-learning layer leads to a consistent reduction in both metrics, with RMSE improvements approaching 10%(reducing from

437 to 394 MW). While the two-layer framework does not fully match the upper-bound model, it recovers a substantial fraction of the lost performance.

This result is particularly significant, as it demonstrates that meaningful gains can be achieved **without access to perfect information** and without retraining the main forecasting model.

Beyond average error metrics, the two-layer framework also exhibits improved behaviour during extreme events. The analysis shows that large forecast errors are disproportionately reduced, indicating that the meta-learner is especially effective at correcting severe input deviations.

This observation is consistent with the ensemble nature of the Adaptive Random Forest, which can capture nonlinear relationships between forecasted and observed weather variables. As a result, the meta-learning layer acts as a dynamic error filter, attenuating the impact of extreme weather forecast inaccuracies.

Comparing the solar and nuclear case studies reveals an important distinction between **structural non-stationarity** and **stochastic uncertainty**. In the solar case, the underlying system remains stable, and the meta-learner operates primarily as a noise corrector. In the nuclear case, the system itself evolves, requiring the meta-learner to infer latent structural variables.

Despite these differences, the same two-layer architecture proves effective in both settings. This highlights the generality of the proposed framework and supports the claim that decoupling prediction and adaptation is a powerful design principle for energy forecasting under non-stationarity.

Overall Discussion and Lessons Learned

Across both case studies, several key insights emerge:

1. **Adaptation without retraining is feasible** for a wide range of non-stationary energy forecasting problems.
2. **Meta-features provide a flexible interface** between adaptive components and fixed forecasters.
3. The framework performs best under **partial but persistent non-stationarity**, while extreme regime shifts remain challenging.
4. The computational efficiency of the approach makes it attractive for large-scale operational deployment.

These findings position the proposed two-layer framework as a practical compromise between static models and fully online learning systems.

CONCLUSIONS

This paper proposes a two-layer meta-learning framework for energy forecasting under non-stationary conditions, addressing both structural concept drift and uncertainty in exogenous inputs. By decoupling prediction from adaptation, the framework enables lightweight online adjustment without retraining the main forecasting model.

Results from real-world case studies demonstrate that the proposed approach consistently improves forecasting robustness across heterogeneous settings. In nuclear power production forecasting, the meta-learning layer effectively tracked capacity-driven structural changes, substantially reducing error relative to static baselines and maintaining meaningful predictive performance under moderate and strong drift. In solar power forecasting, the framework mitigated error propagation from uncertain weather inputs, recovering a significant portion of the performance gap to an upper-bound model trained on observed meteorological data.

An important outcome of this work is the demonstration that adaptation can be localised within a meta-layer, avoiding the computational cost and operational complexity associated with frequent retraining or fully online learning. This makes the proposed framework particularly suitable for large-scale operational forecasting systems.

The framework's limitations become apparent under extreme regime shifts that fundamentally invalidate historical relationships, and the present work is restricted to point forecasting. Future research will therefore focus on extending the approach to probabilistic forecasting, improving robustness to extreme structural breaks, and evaluating its applicability to additional energy system use cases.

Overall, the results indicate that two-layer meta-learning provides a practical and effective strategy for improving energy forecasting performance in evolving power systems.

ACKNOWLEDGEMENTS

This work was supported by the European Union's Horizon 2020 research and innovation program under the Marie Sklodowska-Curie grant agreement No 101150337 – DMOptfES and under the project EGRID (Innovative Tracking, Monitoring and Reduction of Electricity Carbon Emissions) supported by COMPETE2030-FEDER-01073700.

AUTHOR IDENTIFIERS

Marcato, F: 0009-0003-5161-4635
Ruivo de Oliveira, M: 0009-0008-1142-3777
Santecchia, A.: 0000-0002-6089-0776
Francesco Silvestri: 0000-0002-9077-9921
Castro-Amoedo R.: 0000-0001-8287-8295

REFERENCES

1. Santecchia A, Kantor I, Castro-Amoedo R, Maréchal F. Industrial flexibility as demand side response for electrical grid stability. Front. Energy Res. 10: (2022). https://doi.org/10.3389/fenrg.2022.831462
2. Santecchia A, Castro-Amoedo R, Nguyen T, Kantor I, Stadler P, Maréchal F. The critical role of electricity storage for a clean and renewable european economy. Energy Environ. Sci. 16:5350-5370 (2023). https://doi.org/10.1039/d3ee02768f
3. Ali KH, Sigalo M, Das S, Anderlini E, Tahir AA, Abusara M. Reinforcement learning for energy-storage systems in grid-connected microgrids: an investigation of online vs. offline implementation. Energies 14:5688 (2021). https://doi.org/10.3390/en14185688
4. Tsymbal, A. (2004). The Problem of Concept Drift: Definitions and Related Work. Tech. Rep. Department of Computer Science, Trinity College, Dublin.
5. Widmer G, Kubat M. Learning in the presence of concept drift and hidden contexts. Machine Learning 23:69-101 (1996). https://doi.org/10.1023/a:1018046501280
6. Klinkenberg R. Learning drifting concepts: example selection vs. example weighting. IDA 8:281-300 (2004). https://doi.org/10.3233/ida-2004-8305
7. Ikonomovska E, Gama J, Džeroski S. Learning model trees from evolving data streams. Data Min Knowl Disc 23:128-168 (2010). https://doi.org/10.1007/s10618-010-0201-y
8. Gomes HM, Bifet A, Read J, Barddal JP, Enembreck F, Pfharinger B, Holmes G, Abdessalem T. Adaptive random forests for evolving data stream classification. Mach Learn 106:1469-1495 (2017). https://doi.org/10.1007/s10994-017-5642-8
9. Heitor M. Gomes et al. “Adaptive random forests for data stream regression”. In: *ESANN 2018 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* (2018), pp. 267–272.
10. Putatunda S. Supervised learning for streaming data. Practical Machine Learning for Streaming Data with Python :57-96 (2021). https://doi.org/10.1007/978-1-4842-6867-4_3
11. Pinson P. Wind energy: forecasting challenges for its operational management. Statist. Sci. 28: (2013). https://doi.org/10.1214/13-sts445
12. Hong T, Pinson P, Fan S, Zareipour H, Troccoli A, Hyndman RJ. Probabilistic energy forecasting: global energy forecasting competition 2014 and beyond. International Journal of Forecasting 32:896-913 (2016). https://doi.org/10.1016/j.ijforecast.2016.02.001
13. Zhang, Z., Zhang, H., Zhang, Y., & Wang, Y. “Noise-intensification data augmented machine learning for day-ahead wind power forecast”. Energy Conversion and Management 266 (2022), 115803. https://doi.org/10.1016/j.enconman.2022.115803
14. Ren Y, Suganthan PN, Srikanth N. Ensemble methods for wind and solar power forecasting—a state-of-the-art review. Renewable and Sustainable Energy Reviews 50:82-91 (2015). https://doi.org/10.1016/j.rser.2015.04.081
15. Nishanthi M, Pandiyan P. Meta-learning-based imputation in solar energy systems: bridging the gap between missing data and forecasting reliability. J. Electr. Eng. Technol. 21:991-1010 (2025). https://doi.org/10.1007/s42835-025-02420-1
16. Mez L, Okamura L, Weidner H. The ecological modernization capacity of japan and germany. Springer Fachmedien Wiesbaden (2020). https://doi.org/10.1007/978-3-658-27405-4
17. ENTSO-E. “ENTSO-E Transparency Platform”. URL: https://transparency.entsoe.eu/